

# Twenty open problems in quantum control

Kurt Jacobs

Department of Physics, University of Massachusetts at Boston, Boston, MA 02125, USA

The subject of controlling quantum systems is not new, but concepts that have been introduced in the last decade and a half, especially that of coherent feedback, suggest new questions that broaden and deepen the field. Here we provide a concise overview and definition of quantum feedback control, both coherent and measurement-based, and discuss its relationship to “standard” time-dependent control; there is a sense in which the latter subsumes the rest. There are many open questions within quantum control and its subfields, and we highlight and discuss some of them here. These questions are of theoretical as well as practical interest: the answers will help to determine the relative power of the different methods of control, and the limits to our ability to control quantum systems imposed by available resources.

PACS numbers: 03.67.-a, 03.65.Ta, 03.67.Pp, 03.65.Aa

## INTRODUCTION: FEEDBACK AND OPEN-LOOP CONTROL

When one thinks of feedback control, it is likely that the use of measurements will spring to mind. Measurements do not sit easily with quantum theory, and in fact they are somewhat superfluous. Measurements can always be replaced by considering instead the joint evolution of two systems, a primary system and an auxiliary. A measurement process on the primary constitutes the formation of correlations between the two, and feedback is simply the exploitation of these correlations by the joint evolution. In fact it is becoming increasingly clear that the ability to swap between a picture involving explicit measurements and one without measurements is useful. It suggests two alternative ways to physically implement the same control process, one in which measurement results are processed by classical computers, and one that merely involves the application of time-dependent control to a collection of quantum systems. While these two implementations are in a sense equivalent, they are very different from a practical point of view, and both have relative advantages. Further, while coherent feedback subsumes measurement-based feedback, the reverse is not true [6].

To distinguish between the two ways of implementing feedback control we refer to the one that involves explicit measurements as *measurement-based feedback control*, and the one without measurements as *coherent feedback control*. Let us now consider the most general form of feedback control and see how the different categories of control appear within this form. Since one can always subsume measurements into unitary evolution by including an auxiliary system, the most general feedback control scenario need not refer to measurements. We can therefore define this control in the following way. A primary system  $S$  is coupled to an auxiliary system  $C$  (here  $C$  stands for *controller*). Both the system and controller will in general also be coupled to environments (other quantum systems) that inject noise into them, and we

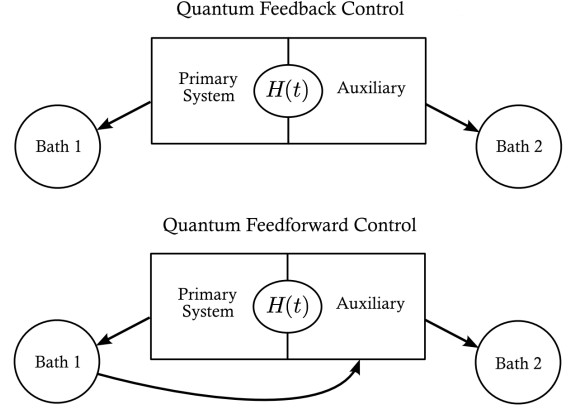


FIG. 1. Depictions of the most general implementation of feedback (top) and feedforward (bottom) control of quantum systems. In the former the auxiliary system, or controller, does not have direct access to the noise driving the system. The direction of the arrows represents the direction in which the flow of information impacts the control problem: In feedback control neither bath provides any useful information to the systems, while in feedforward the auxiliary can extract useful information from bath 1 about the noise it injects into the system.

denote these respectively by  $\mathcal{E}_S$  and  $\mathcal{E}_C$ . The system and auxiliary evolve under a time-dependent Hamiltonian  $H(t)$ , whose purpose is solely to control the evolution of  $S$ . In addition to designing some or all of the time-dependence of  $H$ , we may also have some limited control over the coupling of each of the two systems to their environments. This is the most general form of feedback control, and is depicted in Fig. 1(top).

If the controller  $C$  is a quantum system whose states are not directly accessible to an observer, then the above scenario is coherent feedback control. If some of the states of the auxiliary, or more generally *subspaces* of the auxiliary, are macroscopically distinct and accessible to an observer, then the scenario is that of measurement-based feedback control. In this case the auxiliary must

have within it a means of amplification. That is, the microscopic auxiliary states that become correlated with the microscopic states of the system must be mapped to macroscopically distinct states, whose energy is necessary much larger than that of the microscopic states.

Shortly after posting the first version of this monograph, I had the opportunity to discuss the above definitions of feedback and feedforward with Daniel Lidar, Robert Kosut, Edmond Jonckhere, and Herschel Rabitz. Kosut pointed out that there was more than one definition of feedforward in the classical control community, and further, Jonckhere mentioned that historically some categories of control have been very hard to precisely define. One well-known control theorist at a famous discussion of “adaptive” control at an IEEE meeting remarked that “the only way to know if a control protocol was adaptive was to ask the designer”. It was also suggested that clarifying the role of Markovian vs. non-Markovian baths in these definitions would be useful. In Appendix B I discuss these points in more detail, elucidating the different classical notions of feedback and feedforward. It may be that grey areas between these categories are inevitable, and it may also be important not to be overly zealous in asserting a single definition as superior. Nevertheless it is important to understand the motivation and justification for any definition, and to specify clearly the definition used in a given piece of work. In Appendix A I expand upon the motivations for the loop-free definition of feedback given above.

To introduce some terminology, we note that the purpose, or *objective*, of a control procedure may be to prepare a particular state at a given time (state-preparation), to realize a particular unitary evolution for the system over some time interval, or to have the system evolve in some specified way as a function of time. A specific choice for the auxiliary system and the joint time-dependent Hamiltonian, designed to achieve a specific objective is called a control *protocol*.

There are a few points that are worth noting about the above scenarios, and we pause to discuss these now.

### Everything within open-loop control

If we remove the auxiliary system, or equivalently decide that we wish to control the whole of the joint system,  $J = S \otimes A$ , then the feedback control scenario reduces to that of open-loop control, traditionally referred to simply as *quantum control*, in which the Hamiltonian is varied as function time to achieve an objective. In a sense feedback control is open-loop control in which we wish “only” to control a subsystem of the system to which we apply the control.

### Where is the feedback loop?

Traditionally feedback control involves extracting information from a system, and then using this informa-

tion to determine forces (controls) that are applied to the system. The applied controls are often referred to as being *conditional*, or *conditioned*, on the measurement results. The flow of information to the controller via the measurements, and then back to the system via the control forces, is described as a *loop*. This is why feedback control is also referred to as *closed-loop control*. In the description of feedback control depicted in Fig. 1 we have dropped the loop, and this is something that tends to confuse control theorists. The reason that we do not need to explicitly enforce a loop in which information flows one way is that the auxiliary system has no direct access to the noise that the control protocol is attempting to mitigate. This noise is injected into the system from the environment  $\mathcal{E}_S$ , and the only way that the auxiliary can obtain information about the effects of this noise is to interact with, and thus become correlated with the system. In order to apply forces to the system that can mitigate the noise, the auxiliary must first extract information from the system. The process is therefore necessarily feedback control, purely because the controller does not have direct access to the environment  $\mathcal{E}_S$ .

### Wither feedforward control?

Feedforward control, as defined in classical control theory, is a scenario in which the controller has access to the noise that is affecting the system before it reaches the system. The controller can then determine forces to apply to the system based on this noise and then *feed these forward* to the system to counteract the noise. This is not usually considered for quantum systems because one does not usually have useful access to the environment. But there could be situations in which feedforward is useful for controlling quantum systems, and we depict this situation in Fig. 1(bottom). The only difference between feedback and feedforward is that in the latter the auxiliary system has access to the environment  $\mathcal{E}_S$ .

### Reversible or irreversible coupling?

In measurement-based feedback control one often considers measurements in which information is continually extracted from the system, so called “continuous measurements” [2]. Continuous measurements are realized by coupling the system to an environment, or bath, in which the interaction is irreversible and continuously carries off information about the system. The archetypical example of this is the act of measuring the position of an object by shining a beam of light on it and observing the reflected light. The many modes of the traveling-wave electromagnetic field (the light) are a bath with which the system interacts, and the reflected light continually carries away information.

If instead of measuring the reflected light in the above example, we injected this light into an auxiliary quantum system, and then allowed the auxiliary to act in some way

Meas.-based feedback (MBC)	Coherent feedback (CFC)
noise-free processing requires amplification	controller has noise/damping fast, no amplification

TABLE I. The relative practical merits of CFC and MBC

upon the system, we would have replaced the continuous measurement-based feedback control with a version of coherent feedback control. The information extraction by the controller is mediated in this case by a quantum electromagnetic field. We could similarly apply the feedback forces by coupling the auxiliary to the system by a quantum field, where in this case the field travels from the auxiliary to the system. This scenario explicitly involves information traveling in one direction around a loop. While this puts it explicitly in the form of a classical feedback loop, it necessarily involves a loss of information via the irreversible connection to the fields. In fact, field-mediated coupling, often called a *cascade* connection because of the one-way flow of information [3, 4], is equivalent to a direct (reversible) coupling between the system and auxiliary, with the addition of irreversible coupling of each to separate environments. As a result coherent control via cascade connections is a strict subset of coherent feedback that involves direct coupling as depicted in Fig. 1 (top).

### COHERENT VS. MEASUREMENT-BASED FEEDBACK CONTROL

Techniques for controlling quantum systems via their interaction with another system have been studied and devised for some decades, and these are all ultimately examples of coherent feedback control. Nevertheless many general questions about coherent feedback control are as yet unanswered, as well as many regarding the relationships between coherent feedback control (CFC) and measurement-based control (MBC). First there are two simple practical differences between CFC and MBC that are worth noting. In MBC, since the information is amplified and transformed to classical numbers, this information can be processed by a classical computer. This processing is therefore noise-free and can be rather sophisticated. The disadvantage is that the amplification process, because it must be achieved with very low noise, is technologically challenging and can be slow. The merits of CFC is exactly the converse. Since it requires non measurement or amplification, it can be fast and is technologically simpler, since it only involves coherent manipulation of quantum systems via externally applied classical fields. The flip side is that the controller is now an auxiliary quantum system, which will in general be subjected to noise and dissipation from its environment. Further, the simpler the auxiliary system, the less processing it will be able to do on the information that it obtains from the system.

### Quantifying and comparing dynamical resources

The first problems we consider involve the relationship between CFC and MBC. There is a difference between a reversible coupling, such as the unitary interaction between two microscopic systems, and the irreversible coupling between a system and bath, such as a quantum field. This difference is that in the reversible coupling, so long as the energies of the two systems are bounded, if the systems start in a product state, the evolution of the entanglement generated between the systems is initially second order in time. This is enforced by the unitarity of the evolution. This behavior is distinct from the evolution generated by the interaction with a bath, at least in the Markovian approximation, because the evolution of the entanglement, and the populations of the states of the system is first-order in time. This means that the dynamical resources provided by reversible and irreversible coupling are distinct in some sense.

An important question in any kind of control is how well an objective can be realized given the available resources, and thus the classification of resources is important. When we have at our disposal a unitary interaction between a system and auxiliary, a natural resource to consider is the difference between the maximum and minimum eigenvalues of the interaction Hamiltonian. This quantity is bounded in practice, and bounds the rate at which we can extract information from the system, which in turn bounds the rate at which we can extract entropy and reduce the effects of noise [5]. If we instead have an irreversible coupling to a zero-temperature bath, the natural resource is the size of the damping rate induced by this coupling. It is not clear, however, whether there is a way to compare the eigenvalue bound for unitary interactions to the damping rate of an irreversible coupling. These may be comparable in some theoretically or practically motivated way, or they may be distinct.

**Problem 1:** Can reversible and irreversible couplings be compared quantitatively, and thus be regarded as manifestations of a single resource for quantum control, or do they constitute distinct resources?

Feedback control via continuous measurements is an example of an irreversible Markovian coupling between the primary system and the auxiliary. The efficacy of field-mediated control can therefore be compared directly with that of feedback control via continuous measurements. But a comparison between the latter and coherent feedback control implemented solely with direct (unitary) coupling between the two is only meaningful if unitary and irreversible couplings can be compared.

Some light might be shed on the above question by examining the relationship between field-mediated and direct coherent control. As mentioned above, field mediated control is equivalent to a direct coupling combined with an irreversible bath coupling for both the system

and auxiliary. From this it is clear that control via direct coupling can be compared to continuous measurement-based control if we allow the former to be augmented by coupling the system and auxiliary to appropriate baths.

**Problem 2:** What is the precise relationship between continuous measurement-based feedback, and coherent feedback via direct system-auxiliary coupling given the same dissipative resources? Is the latter superior to the former, and if so, does this remain true without the dissipative resources?

**Problem 3:** Coherent feedback mediated by fields encompasses continuous MBC, except for the processing power of a classical computer. In the former the measurement results must be processed in some way via the evolution of the auxiliary. Determine the size and dynamics required of a quantum auxiliary to approximate various kinds of classical information processing of the measurement record.

### Feedback control via unitary interactions

For coherent feedback control in which the system interacts with the auxiliary via a Hamiltonian whose maximum eigenvalue difference,  $\Delta\lambda \equiv \lambda_{\max} - \lambda_{\min}$  is bounded, it has been shown that coherent feedback control can extract entropy from the system faster than measurement-based feedback under the same bound. This means that CFC will outperform MBC under this bound for a wide range of objectives. To understand why this is, we need to understand more precisely the difference between coherent feedback and measurement-based feedback.

We stated above that all MBC protocols can be reproduced by CFC protocols, but the reverse is not true: MBC protocols are a subset of CFC. To see this we note first that all CFC operations on the primary system, since they involve the interaction with an auxiliary which is then traced out, can be written in the Krauss form

$$\rho(t) = \sum_n A_n \rho(0) A_n^\dagger. \quad (1)$$

for an arbitrary set of operators  $\{A_n\}$  that satisfy  $\sum_n A_n^\dagger A_n = I$ . A general efficient measurement on the primary is also of this form, where each measurement result, labelled by  $n$ , places the system in the state  $\rho_n = A_n \rho A_n^\dagger / \mathcal{N}$  (the normalization is  $\mathcal{N}$ ). We can then perform feedback, which is a unitary operation on the system for each result  $n$ . Averaging over the measurement and feedback, the final state of the system is

$$\rho(t) = \sum_n U_n A_n \rho(0) A_n^\dagger U_n^\dagger. \quad (2)$$

The MBC form looks just as general as the CFC form, but there is a catch. In order that that MBC reduce the entropy of the system, and actually utilize the feedback step (if it didn't it wouldn't be MBC) there must be

more than one measurement outcome, because a single unitary preserves the (von Neumann) entropy. What is shown in [6] is that, under a bound on the Hamiltonian that performs the interaction between the system and auxiliary (an interaction that is also required to make a measure), the fastest way to extract all the entropy from the system using CFC is to follow a geodesic in the joint Hilbert space of the system and auxiliary [6]. The operation on the system that results from this geodesic has only a single Krauss operator:

$$\rho(t)_{\text{geo}} = |\psi\rangle\langle\psi| \rho(0) |\psi\rangle\langle\psi|, \quad (3)$$

where  $|\psi\rangle$  is the final pure state in which the protocol leaves the system. Because of this an MBC protocol cannot follow a geodesic, and must take longer to extract all the entropy.

The result in [6] provides a greater understanding of the relationship between CFC and MBC, and provides a guiding principle for the design of CFC protocols. It suggests that interactions that follow geodesic evolution provide superior performance for noise reduction not only for reversible but also irreversible coupling. Nevertheless...

**Problem 4:** Is it possible to show quantitatively how the restriction to measurement and feedback limits the performance of feedback via irreversible coupling, and how this relates to geodesics?

**Problem 5:** What are the geodesics that correspond to specific control objectives, and how closely can these be approximated by interactions available in mesoscopic systems?

**Problem 6:** Coherent feedback via unitary interactions is inherently a periodic process. What kind of Hamiltonians are required to obtain continuous or quasi-continuous control, and is it possible to obtain simple analytic expressions for steady-state performance in some regimes?

### Limits to quantum control

A practically important and fundamentally interesting question is how well various control objectives can be achieved given specific resources. When controlling mesoscopic systems, a major experimental constraint is the strength, or speed, with which we can interact with the system. The problem with obtaining theoretical limits to feedback control is that to do so we must find optimal protocols. For measurement-based control this is near-impossible, since the equations of motion are both stochastic and nonlinear. But there may be a route to finding upper bounds by exploiting the connection between coherent and measurement-based control. In general, while the equations for coherent feedback control are linear, since they involve only the Schrödinger equation, they are just as complex as the nonlinear equations of MBC. However the problem simplifies if we consider

the regime of good control. Good control implies that the control protocol is able to achieve the objective with a fidelity close to unity. This means that the rate at which it is able to extract entropy from the system is much faster than the rate at which the noise generates entropy, and thus the speed of the interaction with the system is correspondingly fast. Fortunately the regime of good control is of practical importance, and its simplifying feature is that the noise can be treated as a perturbation on the dynamics.

In [5] the authors were able to obtain the optimal (reversibly coupled) coherent feedback protocol for preparing pure states (e.g. for cooling) in the presence of thermal noise, in the regime of good control. The authors assumed an ideal auxiliary system so as to obtain the limits to cooling under a constraint on the speed of the interaction with the system. This shows that it is possible to obtain optimal results for CFC in the regime of good control, and it may therefore be possible to use similar results to obtain bounds on measurement-based control.

In [5] the authors did not solve the most general problem of cooling in the regime of good control, but restricted themselves to the regime in which the interaction with the system does not change its energy eigenvalues appreciably. The reason for this restriction is that it corresponds to the regime in which the control operations do not change the dynamics induced by the thermal noise. The Markovian Redfield master equation that describes this noise depends on the system's energy levels. Thus the authors of [5] solve the problem of optimal CFC for the case in which the master equation is fixed. Relaxing this restriction would provide the answer to the fundamental question of just how well one can prepare pure states given a maximal speed of interaction with the system, taking account of the fact that one can significantly alter the dynamics of the system during the process. While the structure of the more general problem is somewhat different to that solved in [5], the results in [5] suggest that it may be tractable. Further, for a Markovian and therefore memoryless bath, it is easy to determine the Master equation as a function of the control Hamiltonian: the Master equation must immediately adapt to the new Hamiltonian at each point in time. Thus the master equation is merely the Redfield equation implied by the Hamiltonian of the joint system at each time.

**Problem 7:** Determine the optimal protocol for state preparation via a unitary interaction with a system, in the regime of good control, under a bound on the norm of the interaction Hamiltonian, and a bound on the rate at which the Hamiltonian can be changed.

Often there is a practical bound on the interaction strength between two mesoscopic systems, but this does not stop us from coupling a system of interest to multiple auxiliary systems. As we increase the number of auxiliary

systems, the effective coupling rate between the system and the set of auxiliaries increases.

**Problem 8:** Determine the optimal protocol for state preparation via a unitary interaction with the system, in the regime of good control, under a bound on the norm of the interaction Hamiltonian with a given auxiliary, as a function of the number of auxiliary systems.

We mentioned above that it may be possible to bound the performance of continuous MBC by using bounds obtained on CFC, where the CFC is chosen to have equivalent resources (irreversible coupling strength). The reason for this is that CFC subsumes MBC. We also mentioned that reversibly coupled CFC is equivalent to irreversibly coupled CFC with the addition of coupling to specifically chosen baths. Thus by finding optimal solutions for reversibly coupled CFC (as in problems 6 and 7), but with the addition of some irreversible bath coupling, it may be possible to obtain useful bounds on continuous MBC.

**Problem 9:** Determine bounds on the performance of continuous measurement-based feedback control, in the regime of good control, by finding optimal protocols for coherent feedback control using equivalent resources.

### Linear field-mediated feedback control

A subset of CFC problems that is being investigated by a number of groups from the control theory community is that of linear quantum systems coupled (irreversibly) via quantum fields. These systems are also referred to as quantum linear networks. It is well-known that linear systems have a very restricted dynamics, in that they cannot, for example, create non-Gaussian quantum states from Gaussian ones. It is also known that linear networks cannot perform universal quantum computation in an efficient manner. Nevertheless there are other tasks that one would expect that linear networks cannot perform, but proofs have not yet been obtained.

**Problem 10:** Show that linear networks cannot perform quantum error-correction.

Error-correction involves correcting errors on arbitrary states, that is states about which the error correcting process is ignorant. A task that is simpler is protecting known states from the effects of noise and decoherence.

**Problem 11:** Show that linear networks cannot protect non-Gaussian states from linear noise.

So far work on linear feedback networks has been focussed on time-independent networks, in which the Hamiltonians of the systems and the couplings between them do not change with time. The analysis in [5] suggests that including time-dependence in networks can increase their performance for the purposes of control.

**Problem 12:** Quantify the increase in performance provided by including time-dependent external control

in linear feedback networks, and examine how this varies across classes of problems.

We mentioned above that direct (reversible) coupling between systems is more flexible than irreversible (field-mediated) coupling, because the latter comes with bath coupling. This means that a larger class of linear networks can be realized by using direct coupling.

**Problem 13:** Quantify the increase in performance that direct coupling provides over field-mediated coupling, with or without external time-dependent control.

### Imperfect controllers

So far we have considered feedback control as if the auxiliary system were perfect. For measurement based control we can take the processing of the measurement results to be perfect but not necessarily the actions of the controller on the system. For coherent feedback control we cannot even assume that the controller itself, and thus the processing of the information is noiseless. In fact, if the system that we wish to control is noisy, often the auxiliary system will be just as noisy. If this were not the case then why not use the auxiliary in place of the system for that ever task we have in mind?

Sometimes there is a physical justification for why the controller may be different from the system. One example of this is “resolved-sideband” cooling, a simple example of CFC in which a mechanical resonator is cooled by coupling it to an electrical or optical auxiliary resonator. In this case the electrical/optical resonator has such a high frequency that it has essentially zero entropy at the ambient temperature. It is this fact that allows the interaction to transfer entropy out of the mechanical resonator to the auxiliary, cooling the former. But the auxiliary, while it has zero entropy and is in this sense less noisy, still undergoes damping. While this damping is essential to dump the extracted entropy into a reservoir, and thus to obtain steady-state cooling, the analysis in [7] shows that this damping reduces the ability of the control to prepare the purest possible state at a single time (this state is more pure than the best achievable steady-state).

**Problem 14:** Determine optimal control protocols in the regime of good control when the auxiliary system is subject to significant damping.

**Problem 15:** What are the limits to control when the controller is just as noisy as the system? How does this limit scale as the size of the auxiliary is increased in relation to that of the system?

**Problem 16:** How does noise on the control Hamiltonian  $H(t)$  affect the fidelity? Can noise on the auxiliary and noise on the Hamiltonian be managed in the same way?

**Problem 17:** Compare the performance of coherent

and measurement-based control protocols in which the quantum controller is damped and/or noisy. At what level of noise is the advantage provided by coherent feedback eliminated by the noise, and is this important for real applications?

### THE COMPLEXITY OF OPEN-LOOP CONTROL

We now turn to open-loop control, in which there is no auxiliary system, and the task is simply to change the Hamiltonian with time so as to provide a desired evolution. While I have used the word “simply” open-loop control is in fact highly complex. Usually one is faced with a Hamiltonian of the form  $H = \sum_n c_n H_n$ , where the classical parameters  $c_n$  can be varied with time in a more-or-less arbitrary way. The problem of constructing a given unitary  $U$  in some time  $T$  by choosing the  $c_n(t)$  is in general tremendously complex and has no analytic solution. One way to determine the required  $c_n(t)$  is to use numerical search methods. To generate a unitary with dimension  $N$ , given only a fixed number of parameters  $c_n$ , the complexity of the control functions  $c_n(t)$  will in general scale as  $N^2$ , since the number of free parameters in a unitary of size  $N$  has this scaling. Nevertheless, given certain kinds of Hamiltonians  $H_n$ , not all unitaries are equally difficult to produce. In quantum computing, for example, the set of Hamiltonians is two-body interactions between qubits. Those unitaries that can be generated with a sequence of two-body unitaries the length of which scales polynomially with the number of qubits are referred to as fast algorithms (those that scale exponentially are slow).

In the introduction we explained how open-loop control is intimately related to feedback control, since the latter is merely open-loop control on the system and auxiliary. Questions about the difficulty of performing open-loop control therefore have a direct bearing on feedback control. The resources for quantum control are a set of available Hamiltonians  $H_n$  in the joint space of the system and auxiliary. A problem is easy if the complexity of the time-dependent control scales as a polynomial in the logarithm of the size of the system

**Problem 18:** For what control objectives, and what set of control Hamiltonians, is open-loop control easy? Does the set of easy objectives depend in a non-trivial way on the set of control Hamiltonians?

### MISCELLANEOUS PROBLEMS

I promised 20 open problems and so far I have only delivered 18. Here are two more problems to satisfy the title.

**Problem 19:** Continuous-measurement based feedback control is highly complex because modifying the measurement as part of the feedback can increase the performance. Optimal protocols are still not known even

for a single qubit, and progress is hampered by the size of the search space. Explore more complex protocols for small quantum systems than those that have been tried to date to see if they enhance performance, guided by known results to-date. For protocols that do enhance performance, consider whether this enhancement has a bearing on coherent feedback.

**Problem 20:** Are there situations in which feedforward control would be useful in mesoscopic quantum systems? If so explore this paradigm of control with explicit examples and compare it to feedback control.

## ADDENDUM

Quantum control lies at the intersection of information, measurement, and unitary dynamics, and is interesting for this reason. It also has an impact on other fields for which precise control over quantum systems is required, and thus a wide range of potential applications. By painting a broad picture of quantum control, and the complexity and range of open questions in this field, I hope that I have stimulated your interest. With luck you might even feel motivated to investigate some of these questions...

*Acknowledgements:* The author is partially supported by the NSF projects PHY-1005571 and PHY-1212413, and by the ARO MURI grant W911NF-11-1-0268.

## APPENDIX A: REMOVING THE LOOP

In Fig. 2a we depict the standard classical picture of feedback control. The curved arrows represent information flow from the primary system to the auxiliary, and subsequently actions from the auxiliary applied to the system. Time goes clockwise around the loop — the information extraction comes prior to the feedback action. The first important point to note is that the feedback controller can at any time chose to apply an action on the system that does not depend on the information obtained. This action constitutes open-loop control, and thus the latter is necessarily subsumed by feedback control. Thus without affecting our definition of feedback, we can add an arrow going counter-clockwise around the loop from auxiliary to primary. The second important point is that information goes from the system to the auxiliary only if the system causes a change in the auxiliary that is depends in some way on its own state — this generates mutual information. The system can act on the controller in a way that does not depend on its state, but this provides no information to the controller. Further, the actions from the controller on the system could provide the system with information about the controller, and if they do the system could act on the controller in a way that depends on the information it has obtained about the controller. But as far as the feedback control is

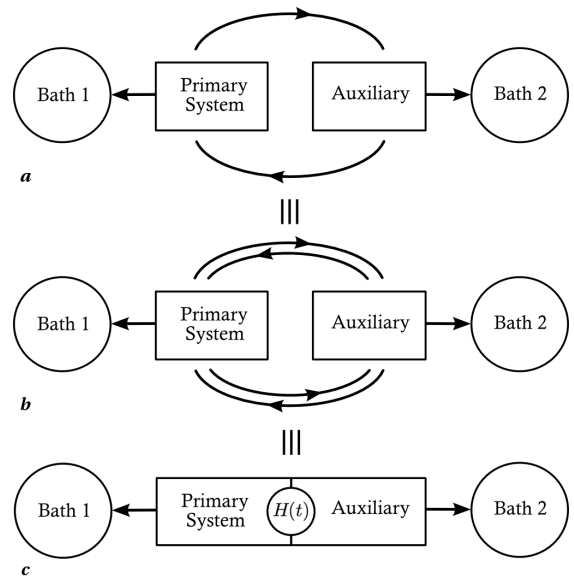


FIG. 2. Here we show the standard definition of classical feedback control (top), along with a situation in which information and conditional actions flow both ways (center), and our definition of feedback control (bottom). In Appendix we explain why these three situations are equivalent.

concerned, whether or not the actions of the system are related to information about the controller, it is *only* the information that they contain about the system that is useful to the controller in controlling the system, any correlations with states of the controller are irrelevant. The ability of the controller to control the system is therefore not improved if we include an arrow going from the system to the controller counter clockwise around the loop. Because of this we are free to include this counter-clockwise arrow within our definition of feedback control; this “extension” does not change the optimal performance of any protocols within the feedback class. In Fig. 2b we include both counter-clockwise arrows, and this “expanded” diagram is equivalent to Fig. 2a. Since the application of forces between two systems with no particular causal direction is simply a time-dependent variation of their joint Hamiltonian, both for classical and for quantum systems, the diagram in Fig. 2c is equivalent to Fig. 2b.

An example of the equivalence of a one-way loop and a two-way Hamiltonian interaction is that of two quantum optical cavity modes connected via one-way fields in a loop. These have the same equations of motion as two oscillators coupled linearly under the rotating-wave approximation, where each oscillator is damped via a quantum-optical master equation. More generally, by coupling quantum systems together via one-way loops we can construct any effective Hamiltonian interaction between the two systems (although this interaction necessarily comes with bath couplings). This is another illustration of why there is no meaningful distinction for

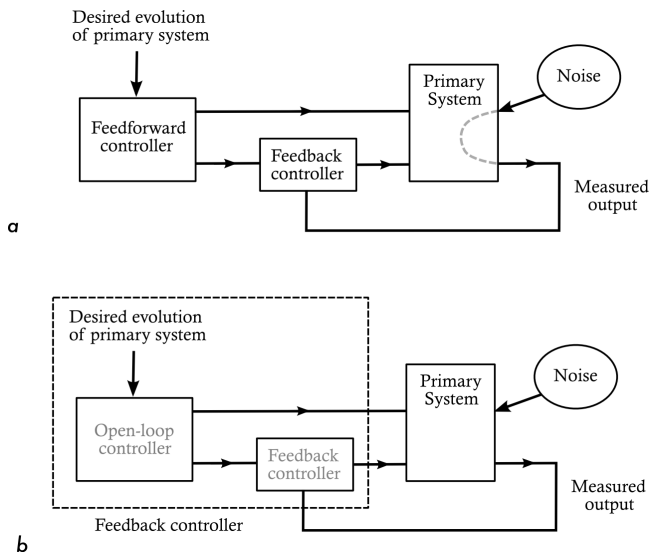


FIG. 3. Top: An alternative classical definition of feedback and feedforward. Bottom: The classical definition of feedback that we use.

our purposes between one-way loops and two-way loops.

## APPENDIX B: TWO DEFINITIONS OF FEEDBACK AND FEEDFORWARD

In Fig. 3a we display an alternative definition of feedforward in classical control systems. We see from the diagram that the box denoting the “feedforward” controller has no access to information from the system, and in our language is therefore referred to as an open-loop controller. Further, in Fig. 3a the feedback controller can be allowed access to the noise realization driving the system at the same time as it effects the system (as indicated by the dashed grey line), but does not have access to the noise *prior* to its effect on the system. It is clear that Fig. 3a is a quite different nomenclature for defining classes of control than the one we use. In Fig. 3b we show how our nomenclature would be used to describe the controller configuration in Fig. 3a. Everything inside the black dashed line is the feedback controller, and we have removed the grey dashed line.

Under our terminology a controller cannot perform feedforward unless it has some access to the noise realization driving the system over and above its interaction with the system. But we have to decide whether or not to include in this definition the situation in which the controller has access to the noise realization at the present time, or to restrict it solely to the case when it has advance knowledge of this noise (that is, *prior* to the time at which it affects the system). To keep our definitions simple, we choose here to define feedback control as the situation in which the controller can interact with the degrees of freedom of the system only, and has no access

to any additional information regarding the noise realization driving the system. It may of course have information about the parameters that characterize the noise, such as the variance(s), correlation function(s), etc. We correspondingly define feedforward control as any situation in which the controller receives information about the noise realization without interacting with the system.

## Markovian vs. non-Markovian baths

The only reason that we include separate baths for the primary and auxiliary in Fig. 1 is to make it clear the auxiliary is unable to extract useful information from the bath that couples to the primary. If the bath’s are Markovian, or more importantly if they are *effectively Markovian*, meaning that the memory time of the bath is shorter than the timescale on which we can modify the joint Hamiltonian, then even when both systems are coupled to the same bath the auxiliary is usually not able to use the bath to extract useful information. Thus in many settings in which there is only a single bath, the situation is that of feedback control, not feedforward (under our nomenclature).

A second question concerns whether it is necessary for bath 1 to be non-Markovian in order to perform feedforward. Under our definition above, the auxiliary is a feedforward controller if it can determine the noise realization driving the system at the current time, without interacting with the system. This is possible if bath 1 is Markovian, and the auxiliary is fed the results of measurements made on it. As an explicit example, if the primary system is a single particle, and the bath induces Markovian dynamics given by the Lindblad operator  $x$  (position), then the effect of the bath is that of a linear random force with a white noise spectrum acting on the particle. If the bath is measured in the right way (in the right basis), then the measurements reveal nothing about the state of the particle, but do reveal the values of the fluctuating force [8], which can then be corrected for by the auxiliary.

- 
- [1] Much less inadequate reference list to come...
  - [2] K. Jacobs and D. A. Steck, *Contemp. Phys.* **47**, 279 (2006).
  - [3] C. W. Gardiner, *Phys. Rev. Lett.* **70**, 2269 (1993).
  - [4] H. J. Carmichael, *Phys. Rev. Lett.* **70**, 2273 (1993).
  - [5] X. Wang, S. Vinjanampathy, F. W. Strauch, and K. Jacobs, *PRL*, in press (2013).
  - [6] K. Jacobs and X. Wang, *Eprint: arXiv:1211.1724* (2013).
  - [7] X. Wang, S. Vinjanampathy, F. W. Strauch, and K. Jacobs, *Phys. Rev. Lett.* **107**, 177204 (2011).
  - [8] A. C. Doherty and Tan, S. M. and Parkins, A. S. and Walls, D. F., *Phys. Rev. A*, **60**, 2380 (1999).